

EvoNova Advisors

The Trust Deficit

Why Finance's Agentic AI Ambitions Are
Stalling at the Threshold of Autonomy

March 2026

EvoNova Insights | Executive Audience

The Controller Who Stopped the Machine

In late 2025, the corporate controller of a Fortune 500 industrial company watched a demonstration that should have been the highlight of her quarter. The company's newly deployed AI agent had just processed 4,200 intercompany reconciliation entries overnight — work that previously consumed three full-time accountants for the better part of a week. The agent had identified 37 exceptions, correctly categorized 34 of them, and drafted resolution recommendations for each. By every efficiency metric, it was a triumph.

Then she asked a question that revealed the chasm at the center of every finance AI deployment in 2026: *"Who approved these entries?"* The answer was no one. The agent had operated within its programmed parameters, but the approval workflow had been designed for human processors. The audit trail showed the agent's logic chain, but not a segregation-of-duties control that any external auditor would accept. Three of the 37 exceptions involved related-party transactions that required controller-level sign-off under company policy. The agent didn't know that — and the system hadn't been designed to ask.

She paused the deployment that afternoon. Not because the technology had failed — by accuracy metrics, it outperformed her team — but because the trust architecture around it did not exist. As she told us:

The AI was faster and more accurate than my team. That wasn't the problem. The problem was that I couldn't explain to our auditors — or our board — how decisions were being made, who was accountable, or what would happen when it got something wrong. Speed without auditability isn't efficiency. It's exposure.

That controller is not an outlier. She is the median. Across the finance and accounting profession, a striking paradox has emerged: the function best suited for agentic AI — data-rich, process-intensive, rules-governed — is the one least willing to let agents act autonomously. Not because the technology doesn't work. Because the infrastructure of trust hasn't been built.

The Paradox of the Perfect Candidate

Consider the data. Finance functions across every major enterprise survey report aggressive plans to adopt agentic AI. According to Deloitte's Q4 2025 CFO Signals Survey, 54% of CFOs at companies with over \$1 billion in revenue rank integrating AI agents as a top transformation priority for 2026. The Savant 2026 Trends Report finds 67% of senior Finance, Tax, and Accounting leaders have pilots underway or implementations in progress. BCG's survey of 280 finance executives at large global companies finds 17% already using GenAI agents, with another 13% planning deployments. The McKinsey Global

Survey on AI reports that 62% of organizations are experimenting with AI agents in at least one function, with finance among the most commonly cited.

Yet beneath the headline adoption figures lies a sobering reality. Only 6% of finance leaders report high maturity levels — what would constitute enterprise-wide, scaled agentic operations. The Wolters Kluwer and CCH Tagetik survey of 392 finance leaders puts it starkly: just 6% currently employ agentic AI, though they project a 600% increase within twelve months. And Gartner predicts that over 40% of agentic AI projects across all functions will be canceled by 2027 due to costs, unclear value, or inadequate risk controls.

The gap between ambition and deployment is not a technology gap. The AI agents available today — from Microsoft’s Copilot for Finance to Oracle’s Fusion AI agents to SAP’s Joule-based workflows — are technically capable of executing multi-step finance workflows: invoice matching, reconciliation, anomaly detection, variance analysis, even aspects of the financial close. The gap is a trust gap. And it manifests at three distinct levels that most organizations have not yet learned to distinguish.

The Trust Architecture: Three Layers That Must Align

The reason most finance organizations remain stuck in what we call **pilot purgatory** — cycling through proofs of concept without reaching production scale — is that they are trying to solve one trust problem when they actually face three. Each layer has different owners, different requirements, and different timelines. Solving one while ignoring the others creates a false sense of readiness that collapses the moment the deployment encounters a real-world stress test: an audit inquiry, a regulatory examination, or a material error (see Exhibit 1).

Exhibit 1:

The Finance Trust Architecture — Three Layers That Determine Whether AI Agents Act or Wait

Three interdependent trust layers, each with distinct requirements and current maturity levels

Trust Layer	Core Question	Key Requirements	Current Maturity	Primary Owner
Technical Trust	Can the agent do it correctly?	Model accuracy, hallucination rates, error handling, data quality integration	Moderate (60-75% accuracy in pilots)	Finance Operations + IT
Process Trust	Can we verify and audit what it did?	Audit trails, explainability, segregation of duties, rollback capability	Low (most lack audit-grade logging)	Internal Audit + Controllers
Institutional Trust	Will regulators, auditors, and boards accept it?	Regulatory compliance, governance frameworks, board-level AI policy, third-party assurance	Very Low (policy gaps in 52% of orgs)	CFO + General Counsel

Source: EvoNova Advisors analysis based on Deloitte Center for Controllershship (2025), Gartner Finance AI Survey (2025), Thomson Reuters Professional Services Report (2025), and Sailpoint AI Agent Adoption Report (2025)

Technical Trust is where most organizations focus their energy — and where the progress is genuinely encouraging. A field study using survey data from 277 accountants and transaction-level records from 79 SMEs reports measurable outcomes: a 9% reallocation of time from routine to higher-value work, a 7.5-day reduction in monthly close time, and a 12% increase in ledger granularity. Consulting case examples describe GenAI assistants replacing 30% of manual variance analysis time at consumer goods companies and AI-enabled invoice scrutiny recovering 4% in contract leakage on multi-billion-dollar spend bases. The PwC AI Agent Survey finds that 66% of adopters cite measurable productivity gains and 57% report cost savings.

But technical capability is necessary and insufficient. Only 14% of CFOs completely trust AI technology to deliver accurate accounting data on its own, according to Wakefield Research. The reason is straightforward: 86% of finance teams have encountered at least one instance of inaccurate or hallucinated data from AI systems. In a profession where a single misclassified journal entry can cascade through financial statements, technical trust requires not just high average accuracy but near-zero catastrophic failure rates — a standard that current models have not yet demonstrated at the level auditors require.

Process Trust is the layer where deployments most commonly fail — and the one least discussed in vendor presentations. Process trust asks: when the agent acts, can we reconstruct exactly what it did, why it did it, and whether the action complied with the organization's control framework? This is not a theoretical concern. The UK Financial Reporting Council's June 2025 landmark guidance on AI in audit explicitly sets documentation and evaluation expectations for AI-enabled tools. The EU AI Act, with high-risk enforcement beginning in August 2026, requires risk management, data quality, logging, traceability, and human oversight for AI systems used in areas like credit scoring and essential financial services.

The governance gap is stark. A Thomson Reuters survey of professional services firms across the US, UK, and Canada found that 52% of respondents believed their organizations had no GenAI policies at work, and 64% reported receiving no GenAI training. A global security survey by Sailpoint found that 82% of organizations using AI agents had policies to secure them, but 80% had observed agents taking unintended actions, and 23% reported agents being tricked into revealing access credentials. In finance — where segregation of duties and approval hierarchies exist for precisely this reason — these findings are alarming.

Institutional Trust is the highest and hardest layer. It asks whether the ecosystem of stakeholders surrounding the finance function — boards, external auditors, regulators, investors — will accept AI agent actions as legitimate inputs to financial reporting and decision-making. The Deloitte Center for Controllershship poll of 3,300 finance professionals provides a telling metric: only 2.7% fully trust AI agents for judgment calls. The vast majority — 59.7% — trust agents only within tightly defined frameworks. The remaining third doesn't trust them at all.

This is not irrational conservatism. The SEC has already taken enforcement action on AI misrepresentation in financial services, settling charges against two investment advisers for misleading

statements about AI use. Regulators are watching. And the institutional trust required for agents to operate in finance cannot be built by the finance function alone. It requires coordinated action across the CFO's office, general counsel, internal audit, and the board — a level of cross-functional alignment that most organizations have not yet attempted.

The Autonomy Gap in Practice

The trust architecture framework explains a pattern visible across every major finance subdomain: a widening gap between where organizations are and where they expect to be. The domains with the strongest technical trust are moving fastest. The domains where institutional trust matters most are stuck (see Exhibit 2).

Exhibit 2:

The Autonomy Gap — Where Finance Functions Actually Stand vs. Where They Expect to Be

Current adoption maturity by finance subdomain, based on cross-referenced survey data (2025–2026)

Finance Subdomain	Current State (Early 2026)	Where Leaders Expect to Be (by 2028)	The Gap	Primary Blocker
Accounts Payable	Workflow agentization (31-68% touchless)	Scaled agentic operations	Moderate	Contract structure / IP rights
FP&A; / Forecasting	Embedded productivity (copilots)	Judgment-layer autonomous forecasting	Large	Model trust & explainability
Reconciliation / Close	Early experimentation (18% piloting)	Workflow agentization	Large	Auditability & error rates
Tax & Compliance	Experimentation (research assistants)	Professional research agents	Very Large	Regulatory uncertainty
Treasury	Embedded productivity (forecasting)	Bounded autonomous execution	Large	Execution risk & governance

Source: EvoNova Advisors analysis based on Savant 2026 Trends Report, BCG Finance Executive Survey (2025), Deloitte CFO Signals Q4 2025, Wolters Kluwer/CCH Tagetik Survey (2025), Gartner Finance AI Forecasts (2025), and PwC Global Treasury Survey (2025)

Accounts payable leads because it sits in a trust sweet spot: high transaction volume, well-defined rules, mature data infrastructure, and limited regulatory sensitivity. Best-in-class AP operations already process invoices at \$2.78 per unit with 3.1-day cycle times, according to Ardent Partners benchmarking. AI-powered invoice matching engines have pushed touchless processing rates from 31% to 68% or higher in documented deployments. The Hackett Group's digital world-class benchmarks show 42% lower process costs and 41% faster close-to-report cycles for organizations that have embraced automation aggressively. AP is where all three trust layers are closest to alignment — which is why it consistently appears as the leading entry point for agentic deployment, cited by 33% of finance leaders in the Savant data.

FP&A and forecasting illustrate the opposite challenge. The technical capabilities are impressive — AI copilots embedded in productivity suites can already draft variance narratives, scenario-model across multiple assumptions, and surface anomalies in real-time. But FP&A outputs feed directly into board presentations, investor communications, and strategic decisions. The institutional trust required for an AI agent to autonomously generate a forecast that a CFO will present to the board is orders of magnitude higher than the trust required for touchless invoice matching. The Maximor Finance AI Adoption Benchmark finds that while 88% of mid-market CFOs use at least one agentic AI tool, 97% insist that human oversight remains critical for ensuring data accuracy. The message is clear: agents can prepare the analysis, but humans must own the conclusion.

Tax and compliance face the most formidable trust barriers. Professional liability, regulatory scrutiny, and the precedent-driven nature of tax interpretation create an environment where even technically superior AI outputs face institutional resistance. Thomson Reuters has positioned its domain-specific agentic assistant for tax and accounting with tool grounding and explainability features specifically designed to manage professional expectations — an architectural acknowledgment that in regulated advisory domains, the agent must not only be right but must prove it was right through a defensible reasoning chain.

The Measurement Trap: Why ROI Alone Won't Close the Trust Deficit

A persistent misconception among finance leaders is that demonstrating ROI from pilot deployments will naturally build the trust needed to scale. It will not. The BCG survey reports a median ROI of 10% from AI and GenAI in finance among those who can quantify it — respectable but not transformative. More revealing: one-third of adopters report ROI below 5%, and only one-fifth report 20% or more. High-impact areas cluster in risk management, forecasting, and exception handling — precisely the domains where institutional trust barriers are highest.

The measurement problem runs deeper than ROI calculation. As MIT Sloan research has cautioned, agentic AI value measurement must separate three distinct categories: gross time reclaimed (hours saved), capacity reallocated (more volume handled with existing staff), and hard-dollar savings (headcount avoided or vendor leakage recovered). Most organizations conflate the three, and 82% of finance leaders in the Savant survey expect no net headcount change from AI in 2026 — suggesting that the primary value proposition is capacity reallocation, not cost elimination. This distinction matters because it changes the trust calculus: reallocation requires trust in the agent's ongoing accuracy; elimination requires trust in the agent's reliability under all conditions.

The organizations breaking through pilot purgatory have recognized that trust is not a byproduct of performance metrics. Trust is an infrastructure investment that must be made *before* performance at scale becomes possible — not after.

Five Moves to Build Trust Architecture Before Your Competitors Do

First, audit your trust layers before you audit your technology. Most finance organizations have conducted technology readiness assessments for AI. Almost none have conducted trust readiness assessments. Map each finance process targeted for agentic deployment against the three trust layers. Where is technical trust established? Where is process trust — auditable logging, segregation-of-duties compliance, explainability — missing? Where does institutional trust require board-level policy or regulatory engagement? The gaps revealed by this exercise will explain why your pilots aren't scaling better than any technology audit ever could.

Second, build audit-grade logging from day one — not as a retrofit. The single most common reason finance AI pilots fail to scale is that they were designed for demonstration, not auditability. Every agent action must produce a log entry that an external auditor can evaluate: what data was accessed, what logic was applied, what decision was made, what alternatives were considered, and what human approval (if any) was obtained. The UK FRC's guidance and the EU AI Act's traceability requirements make this non-negotiable for any organization with European exposure. But even in jurisdictions without explicit requirements, the audit firms are watching. Design for auditability first, and efficiency will follow.

Third, establish a human-in-the-loop escalation framework tiered by risk, not by process. The prevailing approach — requiring human approval for all agent actions in certain processes while granting full autonomy in others — misallocates oversight. A routine invoice match in a \$500 AP transaction and a \$5 million related-party reconciliation entry may run through the same process but carry radically different risk profiles. Design escalation thresholds based on materiality, counterparty sensitivity, regulatory exposure, and precedent — not process category. The Deloitte finding that 59.7% of finance professionals trust agents only within defined frameworks suggests the profession is ready for bounded autonomy. Define the bounds precisely.

Fourth, create a cross-functional AI governance council that reports to the CFO. Trust cannot be built within the finance function alone. It requires alignment between finance operations, internal audit, IT security, legal, and the board. The NIST AI Risk Management Framework and ISO 42001 provide voluntary scaffolding, but they must be translated into company-specific governance that reflects your control environment, your regulatory exposure, and your risk appetite. The 52% of organizations that currently lack any GenAI policy cannot build institutional trust. The CFO is the natural owner of this governance because the finance function sits at the intersection of operational execution, regulatory compliance, and fiduciary accountability.

Fifth, negotiate AI governance into every vendor and outsourcing contract signed after today. As providers embed AI agents into managed service delivery — Oracle's Fusion agents, SAP's Joule workflows, Workday's agent governance layer — the question of who owns the trust architecture

becomes a commercial question. Require contractual provisions for model drift monitoring, bias testing, retraining transparency, and the buyer's right to audit any AI-driven decision that touches financial reporting. Insist on annual commercial rebase windows so that as automation coverage expands, commercial terms adjust. And protect your process IP: as agents learn from your data, ensure you retain rights to any fine-tuned models or process-specific automation scripts built on your operations.

Where This Framework Breaks

The trust architecture model applies most cleanly to large enterprise finance functions with established control environments, external audit relationships, and regulatory obligations. It is less applicable to small and mid-market companies where the controller *is* the governance framework, and where the speed advantages of AI adoption may outweigh formal trust infrastructure investment. Intuit's positioning of AI agents in QuickBooks for small business accounting reflects this reality: the trust model for an owner-operator is personal, not institutional.

The framework also assumes that the technology maturation curve will continue its current trajectory. If large language models develop substantially better reasoning capabilities, lower hallucination rates, and native auditability features within the next twelve to eighteen months, the technical trust layer may resolve itself faster than the institutional layer can adapt. This would create a different problem: technology that outpaces governance, rather than governance that constrains technology. Gartner's projection that 33% of enterprise software applications will include agentic AI by 2028 — up from less than 1% in 2024 — suggests this scenario is not implausible.

Finally, the trust challenge differs by geography. The EU AI Act creates binding obligations with specific timelines. The UK FCA takes a principles-based approach. The SEC focuses on disclosure and enforcement. A global finance function must navigate all three simultaneously, which means the institutional trust layer is necessarily more complex — and more expensive to build — for multinational organizations.

Trust Is the Strategy

The controller who paused her reconciliation deployment did not abandon the technology. She spent eight weeks building what her company now calls its "AI Controls Framework" — a structured protocol that maps every agent action to an approval authority, an audit trail standard, and an escalation threshold. She worked with internal audit to define what agent-generated documentation would satisfy their testing requirements. She worked with legal to draft an AI governance addendum to the company's SOX control documentation. And she worked with the board's audit committee to establish quarterly reporting on AI deployment scope, error rates, and governance maturity.

When she restarted the deployment, it scaled to full production within three months — not because the technology had improved, but because the trust infrastructure made scaling possible. The agent now processes reconciliation entries across 14 entities with a 96.4% accuracy rate. Human reviewers focus exclusively on the exceptions — the 3.6% that actually require judgment. Monthly close cycle time dropped from eleven days to six. And for the first time, external auditors accepted agent-generated work products as audit evidence, citing the documentation framework as the deciding factor.

The broader finance profession stands at precisely this inflection point. The technology is ready. The ambition is unmistakable — 87% of finance leaders predict AI will be extremely or very important to their operations by the end of 2026. The market for AI agents in financial services is projected to grow from \$691 million to \$6.7 billion by 2033. Every major ERP and productivity vendor has embedded agentic capabilities into their platforms.

What remains to be built is not the machine. It is the trust that allows the machine to act.

The organizations that will lead in autonomous finance are not those deploying the most agents. They are those building the deepest trust — the audit trails, the governance frameworks, the escalation protocols, and the institutional confidence that transform a pilot into a production system. In finance, trust is not the obstacle to AI adoption. It is the architecture that makes adoption durable.

About the Author

EvoNova Advisors is a management consulting firm specializing in finance and accounting transformation, intelligent automation, and enterprise operating model design. Our research practice produces original thought leadership on the intersection of technology, process, and commercial strategy.

Contact: insights@evonovaadvisors.com

Sources

McKinsey & Company (2025). The State of AI: Global Survey on AI Agents, June–July 2025.

PwC (2025). AI Agent Survey: 300 Senior Executives, May 2025.

Savant/Wolters Kluwer/CCH Tagetik (2025–2026). 2026 Finance Trends Report and AI Adoption Survey.

Deloitte (2025). Q4 2025 CFO Signals Survey; Center for Controllershship Poll, January 2025.

BCG (2025). Finance Executive Survey on GenAI Agents, March 2025.

Gartner (2025–2026). Finance AI Adoption Survey; Agentic AI Project Forecasts.

NVIDIA (2026). State of AI in Financial Services.

Maximor (2025). Finance AI Adoption Benchmarking Report: 100 Mid-Market CFOs.

Wakefield Research (2025). CFO Trust in AI Survey (cited in Journal of Accountancy, 2026).

The Hackett Group (2026). GBS Key Issues Study; Finance Key Issues: Closing the 5.3% Productivity Gap.

HFS Research (2026). Buying Services-as-Software: How to Make AI Deals That Deliver.

Thomson Reuters (2025). Generative AI in Professional Services Report.

Sailpoint (2025). AI Agent Adoption Report: Enterprise Security Survey.

Ardent Partners (2025). Accounts Payable Metrics That Matter: Benchmarking Study.

UK Financial Reporting Council (2025). Landmark Guidance on AI in Audit, June 2025.

European Union (2024–2027). EU AI Act: Phased Implementation Timeline.

U.S. Securities and Exchange Commission (2024). Enforcement Actions on AI Misrepresentation.

NIST (2023). AI Risk Management Framework. ISO/IEC 42001 AI Management System Standard.

SSRN (2025). Field Study on GenAI in Accounting: 277 Accountants, 79 SMEs (Abstract ID: 5240924).